

Leveraging Uncertainty Quantification to Design Ocean Climate Observing Systems

Nora Loose¹, Patrick Heimbach^{1,2,3}

¹Oden Institute for Computational Engineering and Sciences, The University of Texas at Austin, Austin, TX, USA

²Jackson School of Geosciences, The University of Texas at Austin, Austin, TX, USA

³Institute for Geophysics, The University of Texas at Austin, Austin, TX, USA

Key Points:

- We apply Hessian uncertainty quantification (UQ) to the global ocean state estimate ECCO, and explore its use for observing system design.
- Hessian UQ elucidates oceanic teleconnections that communicate observational constraints over basin-scale distances.
- Going beyond previous adjoint ocean modeling techniques, Hessian UQ rigorously assesses redundancy and optimality of observing systems.

Corresponding author: Nora Loose, nora.loose@colorado.edu

Abstract

Ocean observations are expensive and difficult to collect. Designing effective ocean observing systems therefore warrants deliberate, quantitative strategies. We leverage adjoint modeling and Hessian uncertainty quantification (UQ) within the ECCO (Estimating the Circulation and Climate of the Ocean) framework to explore a new design strategy for ocean climate observing systems. Within this context, an observing system is optimal if it minimizes uncertainty in a set of investigator-defined quantities of interest (QoIs), such as oceanic transports or other key climate indices. We show that Hessian UQ unifies three design concepts. (1) An observing system reduces uncertainty in a target QoI most effectively when it is sensitive to the same dynamical controls as the QoI. The dynamical controls are exposed by the Hessian eigenvector patterns of the model-data misfit function. (2) Orthogonality of the Hessian eigenvectors rigorously accounts for redundancy between distinct members of the observing system. (3) The Hessian eigenvalues determine the overall effectiveness of the observing system, and are controlled by the sensitivity-to-noise ratio of the observational assets (analogous to the statistical signal-to-noise ratio). We illustrate Hessian UQ and its three underlying concepts in a North Atlantic case study. Sea surface temperature observations inform mainly local air-sea fluxes. In contrast, subsurface temperature observations reduce uncertainty over basin-wide scales, and can therefore inform transport QoIs at great distances. This research provides insight into the design of effective observing systems that maximally inform the target QoIs, while being complementary to the existing observational database.

Plain Language Summary

Ocean observing faces multiple challenges: high instrument cost, difficult deployment logistics via ships, harsh environments, and the necessity to sustain observations over long periods of time. Since oceanographers cannot measure the ocean everywhere and at all times, they have to carefully choose the location of their instruments. In an ideal scenario, measurements from a small number of instruments provide maximum information about important ocean metrics, such as poleward ocean heat transport or regional heat content. This paper presents a new method for planning optimal instrument configurations, by combining computer simulations of the global ocean with the mathematics of uncertainty quantification (UQ). As an example, we show that North Atlantic temperature measurements taken below the ocean surface do not only tell us about the ocean properties at the instrument locations themselves, but reduce uncertainty in regions hundreds to thousands of kilometers away. We can therefore use existing ocean observations to extract more information about the ocean than previously appreciated. Our method helps to plan informative observing networks that are complementary to the existing observational database.

1 Introduction

Sustaining long-term ocean observations to develop climate-quality observational records is crucial for understanding the ocean’s role in climate and for evaluating climate model simulations (National Academies of Sciences, Engineering, and Medicine, 2017). Yet, ocean observing faces multiple challenges: complex deployment operations in frequently rough weather (or ice) conditions, limited instrument lifetime due to corrosive and high-pressure environments, and the necessity of adequate spatial and temporal sampling. The high cost and logistical challenges call for deliberate, quantitative approaches. Here, we leverage adjoint modeling and Hessian uncertainty quantification (UQ) within the ECCO (Estimating the Circulation and Climate of the Ocean) framework to explore a new design strategy for ocean climate observing systems. This approach has two distinguishing features, which, taken together, foster collaboration and system co-design within the oceanographic community. First, it gives insights into the physical mecha-

nisms that govern optimal design strategies; and second, it quantitatively assesses redundancy and optimality of an (existing or future) observing system.

To place our technique into context, we briefly recall existing formal approaches to observing system design. Observing System [Simulation] Experiments (OS[S]Es, Fujii et al., 2019) are the most common computational tools in oceanography to support observing system design (e.g., Balmaseda et al., 2007; Gasparin et al., 2019; Griffa et al., 2006; Halliwell et al., 2017). OSEs are limited to evaluate *existing* observing systems, whereas OSSEs can test the skill of proposed *future* observing systems. The design strategy to be tested in an OSSE has to be specified by the investigator. Once a region is targeted for monitoring, the proposed observing system design (to be tested in the OSSE) is typically guided by best available knowledge of both local hydrographic properties and local dynamical balances (Hirschi et al., 2003; Li et al., 2017; Perez et al., 2011). An example are the Atlantic trans-basin mooring arrays OSNAP (Li et al., 2017), RAPID (Hirschi et al., 2003), and SAMBA (Perez et al., 2011), which target monitoring of the meridional overturning circulation. Key components of each design are western and eastern boundary moorings, which allow geostrophic transport estimates across each trans-basin section. Although these local considerations support practical local monitoring, it is possible that similar constraints could be obtained elsewhere, perhaps with an instrument configuration more easily sustained, at reduced cost, or less susceptible to noise. This opportunity arises from the appreciation that observed variability at any given location is rarely a purely instantaneous response to local forcing. Instead, it is the superposition of phenomena originating in distant regions and at distinct times, communicated through the ocean by advection, diffusion and wave propagation (Heimbach et al., 2011).

Exploring the possibility of remote constraint is essential for truly optimal observing system design. Adjoint models have proven valuable for fully mapping the local and remote origins and pathways of variability in targeted quantities, e.g., meridional overturning at given latitudes (Heimbach et al., 2011; Köhl, 2005; Pillar et al., 2016; Smith & Heimbach, 2019). Exploiting the rich information exposed by an adjoint model, a number of adjoint modeling techniques have previously been used to inform ocean observing system design, for example adjoint sensitivity (Heimbach et al., 2011; Masuda et al., 2010), observation sensitivity (Köhl & Stammer, 2004; Moore et al., 2011), and singular vectors (Fujii et al., 2008; Zanna et al., 2012). Despite giving valuable insight into where observations may be useful, none of these latter techniques provide a measure of redundancy versus complementarity, nor of optimality of an observing system. These obstacles are overcome by Hessian UQ: an adjoint-based technique embedded in a variational data assimilation system. The Hessian matrix (composed of second derivatives) of the cost function J captures the curvature of J with respect to the control variables, and allows one to calculate how much uncertainty is reduced with any changes applied to the observing system (Thacker, 1989). In contrast to the previous adjoint modeling techniques named above, Hessian UQ accounts for data redundancy. It also provides a measure of optimality: the more uncertainty an observing system reduces in a defined target quantity (on a scale of 0% to 100%), the closer it is considered to being optimal for the defined target.

Hessian UQ has been routinely applied in numerical weather prediction (NWP, Leutbecher, 2003) and, more broadly, in computational science and engineering (CSE, Bui-Thanh et al., 2012), but it has only seen limited use in the oceanographic community. Previous studies have applied Hessian UQ after severely reducing the dimension of the space of uncertain parameters in an ad-hoc manner (Kaminski et al., 2015, 2018), or in the dual form of ‘representers’ (Bennett, 1985; Moore et al., 2017; Zhang et al., 2010). These examples have focused on regional ocean settings and on daily to monthly time scales. In this study, we take a step toward fully exploiting Hessian UQ to design ocean observing systems that are targeted at climate monitoring in a global context. To this aim, we apply Hessian UQ within the global ocean state estimation framework of the Es-

timating the Circulation and Climate of the Ocean (ECCO) consortium (Heimbach et al., 2019), and elucidate oceanic teleconnections that communicate observational constraints over basin-scale distances and monthly to interannual time scales.

In ocean climate research, the goal of an observing system is usually to accurately estimate certain quantities of interest (QoIs): forecasts or climate indices that are difficult or impossible to observe directly. Examples of QoIs include transports across certain oceanographic passages, ocean heat content near the polar ice sheets, regional sea level anomalies, or future sea-ice extent. We therefore focus on the information that an observing system contains about a given QoI, here referred to as the observing system’s ‘proxy potential’ for the QoI on a scale of 0% to 100% (Loose et al., 2020). Proxy potential is defined by way of Hessian UQ, as the reduction in QoI uncertainty that would be achieved if the observing system was added to the ocean state estimate. Importantly, proxy potential can be assessed not only for existing but also for future observing systems, because it does not require the actual measurement values of the observations (only their locations, times, types, and uncertainties).

Loose et al. (2020) provided interpretations of Hessian UQ and proxy potential for idealized cases, in which an observing ‘system’ consists of only a single and noise-free observation. Then, the observation’s proxy potential for a QoI reflects the degree to which adjustment mechanisms are shared between the observation and QoI. In this simple case, proxy potential can be understood as the dynamical analogue of statistical correlation (squared) between observation and QoI, with the important distinction that proxy potential accounts only for covariability that has dynamical underpinnings. The goal of this study is to leverage Hessian UQ to generalize the notion of proxy potential introduced by Loose et al. (2020) in three important ways (section 2): first, by extending this concept from a single observational asset to full observing systems; second, by quantifying observational redundancy versus complementarity; and third, by accounting for observational noise.

We illustrate the concepts of Hessian UQ and proxy potential in a North Atlantic case study (section 3). To provide a clear understanding of Hessian UQ, our case study focuses on observing systems that are comprised of only a few observations. We then discuss how our approach and the dynamical insights obtained generalize to the design of full-fledged observing systems, including thousands to millions of observations (section 4).

2 Uncertainty Quantification and Proxy Potential

2.1 Ocean state estimation

Ocean state estimation optimally fits an ocean general circulation model (GCM) to the available observations in a dynamically and kinematically consistent way. For this, one solves an inverse problem: given an observing system (gray box, Fig. 1), one adjusts the control vector $\mathbf{u} = (u_1, \dots, u_N)^T$ (green box, Fig. 1), such as to minimize the scalar cost function

$$J(\mathbf{u}) = \underbrace{\frac{1}{2}(\mathbf{y} - \mathbf{Obs}(\mathbf{u}))^T \mathbf{R}^{-1} (\mathbf{y} - \mathbf{Obs}(\mathbf{u}))}_{J_{\text{misfit}}(\mathbf{u})} + \underbrace{\frac{1}{2}(\mathbf{u} - \mathbf{u}_0)^T \mathbf{B}^{-1} (\mathbf{u} - \mathbf{u}_0)}_{J_{\text{prior}}(\mathbf{u})}. \quad (1)$$

The control variables u_1, \dots, u_N (i.e., the elements of the control vector \mathbf{u}) are the uncertain input variables of the model, and consist not only of initial conditions (as common in NWP), but also of atmospheric forcing variables and uncertain model parameters (green box, Fig. 1). The function $J_{\text{misfit}}(\mathbf{u})$ measures the misfit between the vector of actual observations, $\mathbf{y} = (y_1, \dots, y_M)^T$ (gray box, Fig. 1), and the vector of simulated observations, $\mathbf{Obs}(\mathbf{u}) = (\text{Obs}_1(\mathbf{u}), \dots, \text{Obs}_M(\mathbf{u}))^T$ (pink box, Fig. 1), given the input variables \mathbf{u} . The function $J_{\text{prior}}(\mathbf{u})$ penalizes deviations from a first-guess \mathbf{u}_0 of uncertain inputs. The $M \times M$ matrix \mathbf{R} and $N \times N$ matrix \mathbf{B} are chosen error covari-




Figure 1. Workflow for Hessian uncertainty quantification (UQ) in ocean state estimation. Starting from an observing system (gray box), inverse uncertainty propagation along path (UQ1) reduces the uncertainty in the control variables (green box), see section 2.2. A subsequent forward uncertainty propagation along path (UQ2) reduces the uncertainty in a chosen quantity of interest (QoI, purple box), see section 2.3. Green and black arrows indicate propagation of prior and posterior uncertainty, respectively. The degree to which the observing system reduces uncertainty in the QoI, via a composite uncertainty propagation along paths (UQ1) and (UQ2), is referred to as the observing system’s proxy potential for the QoI (section 2.4).

ances, spelling out the assumption that observational noise and prior uncertainties follow the Gaussian distributions $\mathcal{N}(\mathbf{0}, \mathbf{R})$ and $\mathcal{N}(\mathbf{u}_0, \mathbf{B})$, respectively (Tarantola, 2005).

The solution of the inverse problem is the minimizer of the cost function, $\mathbf{u}_{\min} = \min_{\mathbf{u}} J$; that is, a choice of control variables. The ocean state estimate itself is obtained by running the GCM with inputs \mathbf{u}_{\min} .

2.2 Inverse uncertainty propagation

To quantify uncertainties in the solution \mathbf{u}_{\min} of the inverse problem, one propagates observational information and uncertainty along path (UQ1) (Fig. 1). This inverse uncertainty propagation results in the posterior probability distribution of the control variables, given the observations. In practice, it is not feasible to compute the full posterior probability distribution, nor to map this distribution onto the full ocean state space. We therefore need to appeal to approximation methods.

The posterior probability distribution of the control variables can be approximated by the Gaussian $\mathcal{N}(\mathbf{u}_{\min}, \mathbf{P})$, with $N \times N$ covariance matrix \mathbf{P} equal to

$$\mathbf{P} = \mathbf{B} - \sum_{i=1}^{M'} \frac{\lambda_i}{\lambda_i + 1} \left(\mathbf{B}^{1/2} \mathbf{v}_i \right) \left(\mathbf{B}^{1/2} \mathbf{v}_i \right)^T. \quad (2)$$

Here, $\{\mathbf{v}_i, \lambda_i\}_{i=1}^{M'}$ is the set of orthonormal eigenvectors \mathbf{v}_i with associated non-zero eigenvalues $\lambda_1 \geq \dots \geq \lambda_{M'} > 0$ of the misfit Hessian:

$$\mathbf{H}_{\text{misfit}} = \mathbf{B}^{T/2} \mathbf{A}^T \mathbf{R}^{-1} \mathbf{A} \mathbf{B}^{1/2} = \sum_{i=1}^{M'} \lambda_i \mathbf{v}_i \mathbf{v}_i^T. \quad (3)$$

In eq. (3), the entries of the $M \times N$ matrix \mathbf{A} are the sensitivities

$$\mathbf{A}_{i,j} = \frac{\partial(\text{Obs}_i)}{\partial u_j}, \quad (4)$$

evaluated at \mathbf{u}_{\min} . Furthermore, $\mathbf{B}^{1/2}$ denotes an $N \times N$ matrix which has an inverse, $\mathbf{B}^{-1/2}$, and satisfies $\mathbf{B}^{1/2} \mathbf{B}^{T/2} = \mathbf{B}$ (where $\mathbf{B}^{T/2}$ is the transpose of $\mathbf{B}^{1/2}$). The $N \times N$ matrix $\mathbf{H}_{\text{misfit}}$ is the linearized Hessian (or Gauss-Newton Hessian, Chen, 2011) of the rescaled model-data misfit term, $J_{\text{misfit}}(\tilde{\mathbf{u}})$ (eq. (1)). The rescaling is performed through the change of variables $\tilde{\mathbf{u}} = \mathbf{B}^{-1/2} \mathbf{u}$, and can be thought of as a nondimensionalization if \mathbf{B} is diagonal. In summary, eq. (2) phrases the posterior uncertainty \mathbf{P} as the prior uncertainty \mathbf{B} , reduced by any information $\{\mathbf{v}_i, \lambda_i\}$ obtained from the observations. Expression (2) has been known and used in the NWP and CSE communities for many years (see, e.g., Bui-Thanh et al., 2012; Leutbecher, 2003). A self-contained derivation is relegated to the supporting information (Text S1).

Next, we inspect the set $\{\mathbf{v}_i, \lambda_i\}$ in more detail as it fully characterizes the information obtained from the observations. The eigenvectors $\{\mathbf{v}_i\}_{i=1}^{M'}$ of the misfit Hessian (eq. (3)) are the data-informed directions within the control space. Along a data-informed direction \mathbf{v}_i , the function $J_{\text{misfit}}(\tilde{\mathbf{u}})$ has curvature $\lambda_i > 0$ (Fig. 2(a)). The eigenvalue λ_i captures the strength of the data constraint imposed on the control direction \mathbf{v}_i , with large λ_i corresponding to a strong observational constraint. The control directions along which $J_{\text{misfit}}(\tilde{\mathbf{u}})$ is not curved are not informed by the observations (Fig. 2(b)). Note that $M' \leq \min(M, N)$; that is, the number of independent data-informed directions, M' , is less than or equal to the number of observations, M , and the number of control variables, N .

If an observing system consists of only a single observation ($M = 1$) with simulated counterpart $\text{Obs}_1(\mathbf{u}) = \text{Obs}(\mathbf{u})$ and observational noise variance $\mathbf{R} = \varepsilon^2 > 0$,




Figure 2. (a),(b) Curvature of the rescaled model-data misfit function, $J_{\text{misfit}}(\tilde{\mathbf{u}})$, at the cost function minimizer $\tilde{\mathbf{u}}_{\min}$, along two directions in the control space: (a) the data-informed direction \mathbf{v}_i (eq. (3)) and (b) a non-informed direction. (c) The direction of interest, \mathbf{q} (eq. (10)), orthogonally decomposed into $\mathbf{q} = \mathbf{q}_{\text{obs}} + \mathbf{q}_{\text{null}}$. The data-informed component, \mathbf{q}_{obs} , is the projection of \mathbf{q} onto the data-informed subspace. The component \mathbf{q}_{null} lies in the nullspace, i.e., the subspace that is not informed by the data. Parts of the unit sphere of the control space are displayed in black, and \mathbf{q} has unit length. The larger the radius of the orange dashed circle, defined by the length of \mathbf{q}_{obs} , the higher the dynamical proxy potential of the considered observing system for the QoI.

the misfit Hessian (eq. (3)) simplifies to $\mathbf{H}_{\text{misfit}} = \lambda_1 \mathbf{v}_1 \mathbf{v}_1^T$, with

$$\mathbf{v}_1 = (\sigma_{\text{Obs}}^{\mathbf{B}})^{-1} \left[\mathbf{B}^{T/2} \nabla_{\mathbf{u}} \text{Obs} \right] \in \mathbb{R}^N, \quad \lambda_1 = \frac{(\sigma_{\text{Obs}}^{\mathbf{B}})^2}{\varepsilon^2} > 0. \quad (5)$$

Here, \mathbb{R}^N refers to the N -dimensional vector space of real numbers, and we denote $\nabla_{\mathbf{u}} \text{Obs} = (\partial(\text{Obs})/\partial u_1, \dots, \partial(\text{Obs})/\partial u_N)^T$, evaluated at \mathbf{u}_{min} , and

$$\sigma_{\text{Obs}}^{\mathbf{B}} = \left\| \mathbf{B}^{T/2} \nabla_{\mathbf{u}} \text{Obs} \right\| > 0. \quad (6)$$

Put differently, the only data-informed direction is spanned by the prior-weighted sensitivity vector $\mathbf{B}^{T/2} \nabla_{\mathbf{u}} \text{Obs}$ (eq. (5)), where ‘prior-weighting’ is through multiplication by $\mathbf{B}^{T/2}$. Similarly, for an observing system with more than one observation ($M > 1$), the data-informed subspace of the control space is spanned by the M prior-weighted sensitivity vectors $\mathbf{B}^{T/2} \nabla_{\mathbf{u}} \text{Obs}_1, \dots, \mathbf{B}^{T/2} \nabla_{\mathbf{u}} \text{Obs}_M$. To obtain the eigenvectors of the misfit Hessian, one has to orthonormalize and rotate these M vectors within the data-informed subspace (Appendix A). In particular, the eigenvectors of our misfit Hessian – which contains the *second* derivatives of J_{misfit} – are fully determined by *first* (rather than second) derivatives of the observed quantities, i.e., by $\nabla_{\mathbf{u}} \text{Obs}_i$.

For $M = 1$, the observational noise, ε^2 , appears in the denominator of λ_1 (eq. (5)). In particular, for vanishing ε^2 , the eigenvalue λ_1 tends to infinity. This fact generalizes to the case $M > 1$: in the limit of vanishing observational noise ($\mathbf{R} \searrow 0$), the eigenvalues λ_i of the misfit Hessian (eq. (3)) tend to infinity,

$$\lambda_i \nearrow \infty. \quad (7)$$

That is, $J_{\text{misfit}}(\tilde{\mathbf{u}})$ becomes increasingly curved (Fig. 2(a)) and the constraint by the observations increasingly strong.

2.3 Forward uncertainty propagation

To assess the observational constraints on a QoI, the inverse uncertainty propagation along path (UQ1) has to be followed by a forward uncertainty propagation along path (UQ2) (Fig. 1). In other words, we quantify how the uncertainty reduction in the controls, due to the new observational information, reduces uncertainty in the QoI, a diagnostic of the model evaluated at \mathbf{u}_{min} . Forward propagation of prior uncertainties (\mathbf{B} , dotted green arrow) and posterior uncertainties (\mathbf{P} , dotted black arrow) along path (UQ2) results in the prior and posterior QoI variances (see Isaac et al., 2015, or Text S2 in supporting information):

$$(\sigma_{\text{QoI}}^{\mathbf{C}})^2 = (\nabla_{\mathbf{u}} \text{QoI})^T \mathbf{C} (\nabla_{\mathbf{u}} \text{QoI}) = \left\| \mathbf{C}^{1/2} \nabla_{\mathbf{u}} \text{QoI} \right\|^2, \quad \mathbf{C} \in \{\mathbf{B}, \mathbf{P}\}. \quad (8)$$

We infer the prior-to-posterior reduction in QoI uncertainty, relative to the prior uncertainty:

$$\Delta \sigma_{\text{QoI}}^2 = \frac{(\sigma_{\text{QoI}}^{\mathbf{B}})^2 - (\sigma_{\text{QoI}}^{\mathbf{P}})^2}{(\sigma_{\text{QoI}}^{\mathbf{B}})^2} = \sum_{i=1}^{M'} \frac{\lambda_i}{\lambda_i + 1} (\mathbf{q} \bullet \mathbf{v}_i)^2 \in [0, 1]. \quad (9)$$

The second equality in eq. (9) holds by virtue of eqs. (8) and (2). Here, $\{(\mathbf{v}_i, \lambda_i)\}_{i=1}^{M'}$ are the eigenvectors and eigenvalues of the misfit Hessian (eq. (3)), \bullet denotes the ‘dot’ (or Euclidean inner) product between two vectors in \mathbb{R}^N , and

$$\mathbf{q} = (\sigma_{\text{QoI}}^{\mathbf{B}})^{-1} \left[\mathbf{B}^{T/2} \nabla_{\mathbf{u}} \text{QoI} \right] \in \mathbb{R}^N. \quad (10)$$

The unit vector \mathbf{q} is of key interest: it is the direction within the control space to be constrained in order to inform the QoI. It can be written as the orthogonal decomposition $\mathbf{q} = \mathbf{q}_{\text{obs}} + \mathbf{q}_{\text{null}}$ (Fig. 2(c)). \mathbf{q}_{obs} is the component that lies in the data-informed subspace, given by the projection $\mathbf{q}_{\text{obs}} = \sum_{i=1}^{M'} (\mathbf{q} \bullet \mathbf{v}_i) \mathbf{v}_i$. The component \mathbf{q}_{null} lies in the orthogonal complement of the data-informed subspace: the null space, i.e., the subspace that is not informed by the data. Uncertainty is only reduced along the data-informed component, \mathbf{q}_{obs} , not along the nullspace component, \mathbf{q}_{null} .

2.4 Dynamical and effective proxy potential

Relative reduction in QoI uncertainty, $\Delta\sigma_{\text{QoI}}^2$ (eq. (9)), rigorously quantifies the dynamical constraints of an observing system (gray box, Fig. 1) on a QoI (purple box, Fig. 1), as the result of the composite uncertainty propagation along paths (UQ1) and (UQ2). We refer to $\Delta\sigma_{\text{QoI}}^2$ as the *proxy potential* of the observing system for the QoI (Loose et al., 2020). Building on eq. (9), we distinguish between *dynamical* proxy potential

$$\text{DPP}(\text{Obs}_1, \dots, \text{Obs}_M; \text{QoI}) = \sum_{i=1}^{M'} (\mathbf{q} \bullet \mathbf{v}_i)^2 \in [0, 1] \quad (11)$$

and *effective* proxy potential

$$\text{EPP}(\text{Obs}_1, \dots, \text{Obs}_M; \text{QoI}) = \sum_{i=1}^{M'} \frac{\lambda_i}{\lambda_i + 1} (\mathbf{q} \bullet \mathbf{v}_i)^2 \in [0, 1] \quad (12)$$

of the examined observing system for the QoI. Recall that $M' \leq M$ is the number of independent data constraints, characterized by the eigenvectors and eigenvalues $\{\mathbf{v}_i, \lambda_i\}_{i=1}^{M'}$ of the misfit Hessian (eq. (3)). Geometrically, DPP is equal to the squared length of \mathbf{q}_{obs} , the data-informed component of \mathbf{q} in control space (Fig. 2(c)). Note that EPP is smaller than DPP, because all factors $\eta_i = \lambda_i/(\lambda_i + 1)$ are smaller than 1. For vanishing observational noise, EPP approaches DPP, since all eigenvalues λ_i tend to infinity (eq. (7)), and consequently all factors $\eta_i = \lambda_i/(\lambda_i + 1)$ tend to 1 (see also Appendix B and Fig. 7). The bounds for DPP and EPP correspond to the cases for which the observing system provides *no* constraint (EPP = DPP = 0), and for which it serves as a *perfect* proxy for the QoI, in the case of noise-free observations (DPP = 1) and noisy observations (EPP \nearrow 1).

If the observing system under consideration consists of only a single observation ($M = 1$), eq. (11) simplifies to $\text{DPP}(\text{Obs}_1; \text{QoI}) = (\mathbf{q} \bullet \mathbf{v}_1)^2$, which coincides with the definition of dynamical proxy potential in Loose et al. (2020, eq. (4) therein).

3 Application to the North Atlantic

We illustrate the concepts of Hessian UQ and proxy potential in a North Atlantic case study. Section 3.1 describes our experimental setup, including our choice of QoI and observations. We then assess proxy potential of the observations for the QoI, for the cases of noise-free observations (DPP, section 3.2) and noisy observations (EPP, section 3.3).

3.1 Experimental Setup

Our experimental setup coincides with the one described in section 3.1 of Loose et al. (2020) and is embedded in the ECCO version 4, release 2 (ECCOV4r2, Forget et al., 2015) state estimation framework. We use the Massachusetts Institute of Technology general circulation model (Marshall et al., 1997; Adcroft et al., 2018), in a global configuration, at a nominal horizontal resolution of 1° , and with 50 vertical levels. The linear sensitivities of the QoI and observed quantities to all control variables (eqs. (5), (10), and Appendix A) are computed using the respective adjoint models generated through algorithmic differentiation with the commercial tool Transformation of Algorithms in Fortran (TAF, Giering & Kaminski, 2003).

Our QoI is heat transport across the Iceland-Scotland ridge (black line, Fig. 3), denoted by HT_{ISR} . We study four different hypothetical temperature observations in the North Atlantic, located inside the yellow dots in Fig. 3, and labeled by θ^A , θ^B , θ^C , θ^D . Observations θ^A and θ^C are in the Irminger Sea at ($40^\circ\text{W}, 60^\circ\text{N}$), observation θ^B off the Portuguese coast at ($12^\circ\text{W}, 41^\circ\text{N}$), and θ^D in Denmark Strait at ($28^\circ\text{W}, 66^\circ\text{N}$). θ^A ,




Figure 3. Overview map of the case study in this work (modified from Fig. 2 in Loose et al., 2020). The QoI is heat transport across the Iceland-Scotland ridge (black line), denoted by HT_{ISR} . The temperature observations θ^A , θ^B , θ^C , θ^D are located inside the yellow dots. θ^A , θ^B , θ^D are subsurface (at 300 m depth), θ^C at the sea surface. The arrows show approximate pathways of near-surface currents carrying warm Atlantic waters (orange) and cold Arctic waters (purple): NAC = North Atlantic Current; NwAC = Norwegian Atlantic Current; IC = Irminger Current.

Table 1. Control Variables and Prior Uncertainties in our Case Study.

m	Forcing $F_m(i, j)$	Symbol	Time average	# Variables	ΔF_m
1	Net upward surface heat flux	$Q_{\text{net}, \uparrow}$	five years	$6 \cdot 10^5$	50 W/m^2
2	Net surface freshwater flux	EPR	five years	$6 \cdot 10^5$	$5 \cdot 10^{-8} \text{ m/s}$
3	Zonal wind stress	τ_x	five years	$6 \cdot 10^5$	0.05 N/m^2
4	Meridional wind stress	τ_y	five years	$6 \cdot 10^5$	0.05 N/m^2

θ^B , and θ^D are subsurface observations, situated at 300 m depth, whereas θ^C is a surface observation. Exact definitions for the model calculations of HT_{ISR} and θ^* can be found in Loose et al. (2020). We quantify dynamical and effective proxy potential of the five-year mean of the observations for the five-year mean of our QoI, for zero lag. Sensitivities of the QoI (eq. (10)) and observations (eq. (5)) to the control variables are computed from the final five years (2007-2011) of the ECCOv4r2 state estimate.

As in Loose et al. (2020), the control variables are comprised of the time-mean, spatially-varying forcing fields $Q_{\text{net}, \uparrow}$, EPR , τ_x , and τ_y (Table 1). These two-dimensional forcing fields, $F_m(i, j)$, $m = 1, \dots, 4$, are flattened and concatenated into a long control vector, $\mathbf{u} = (u_1, \dots, u_N)^T$. The length of the control vector, N , is $\mathcal{O}(10^6)$, equal to 4 times the number of model surface cells covering the global ocean (next to last column, Table 1). We assign each of the four forcing fields, $F_m(i, j)$, a spatially constant prior standard deviation, ΔF_m (last column, Table 1; see also Loose et al., 2020). Moreover, prior cross-correlations are set to zero, corresponding to the assumption that the decorrelation length in the surface forcing is less than the grid scale ($\sim 1^\circ$). These choices result in a diagonal $N \times N$ prior covariance matrix, \mathbf{B} , whose diagonal is populated by the $(\Delta F_m)^2$ assigned in Table 1. The rationale for our simplified choice of control variables and prior covariance matrix is that it allows a more concise presentation of the results, facilitating a clear understanding of the methodology presented in section 2 and its underlying dynamical and mathematical concepts.

3.2 Noise-free Observations

We begin by exploring the DPP of the temperature observations θ^* , $\star = A, B, C, D$, for our QoI, HT_{ISR} . DPP is equal to the relative uncertainty reduction in HT_{ISR} that is achieved when adding θ^* to the underlying state estimation framework, in the limit of vanishing observational noise (eq. (11)). We first study the DPP of each observation individually (section 3.2.1), and then the DPP of observing systems that are formed by multiple θ^* (section 3.2.2).

3.2.1 Degree of Shared Adjustment Mechanisms

The DPP of each observation θ^* for HT_{ISR} quantifies the degree to which the direction of interest, \mathbf{q} , projects onto the θ^* -informed subspace of the control space (eq. (11) and Fig. 2(c)). We denote the (one-dimensional) θ^* -informed direction by \mathbf{v}^* . As an example, Figs. 4(a)-(e) show \mathbf{q} (eq. (10)) and \mathbf{v}^* (eq. (5)), restricted to their τ_y (meridional wind stress) components:

$$\mathbf{q}_{|\tau_y} = (\sigma_{\text{HT}}^{\mathbf{B}})^{-1} \left[\frac{\partial(\text{HT}_{\text{ISR}})}{\partial \tau_y(i, j)} \Delta \tau_y \right], \quad \mathbf{v}^*_{|\tau_y} = (\sigma_{\star}^{\mathbf{B}})^{-1} \left[\frac{\partial \theta^*}{\partial \tau_y(i, j)} \Delta \tau_y \right]. \quad (13)$$

Here, the normalization factors $\sigma_{\text{HT}}^{\mathbf{B}} = 6 \text{ TW}$ and $\sigma_{\star}^{\mathbf{B}}$, whose values are reported in Table 2, are prior uncertainties, computed according to eqs. (8) and (6), respectively. In eq. (13), prior-weighting, i.e., multiplication by $\mathbf{B}^{T/2}$, has simplified to an element-wise




Figure 4. (a)-(e): Sensitivities of five-year mean (a) heat transport across the Iceland-Scotland ridge (HT_{ISR}), (b) subsurface temperature in the Irminger Sea (θ^A), (c) subsurface temperature off the Portuguese coast (θ^B), (d) surface temperature in the Irminger Sea (θ^C), (e) subsurface temperature in Denmark Strait (θ^D) to changes in five-year mean meridional wind stress τ_y . The solid black-yellow line in (a) delineates the ISR, and the yellow dots in (b)-(e) mark the respective locations of the temperature observation. The sensitivities in (a)-(e) are weighted and normalized (thus unitless), and are the τ_y restriction of the vectors \mathbf{q} (in (a)) and \mathbf{v}^* (in (b)-(e)), see eq. (13). As an example, the insets in (a) and (b) show two arbitrarily chosen vector entries, respectively: q_l, q_m (entries of \mathbf{q}), and v_l, v_m (entries of \mathbf{v}). (f)-(j): Relative contributions of the control fields $Q_{net,\uparrow}$, EPR , τ_x and τ_y (Table 1) to the magnitude of the vectors (f) \mathbf{q} , (g) \mathbf{v}^A , (h) \mathbf{v}^B , (i) \mathbf{v}^C , (j) \mathbf{v}^D , see eq. (14). The yellow labels show the DPP of θ^* for HT_{ISR} (eq. (15)).

multiplication of the sensitivity vectors by the constant $\Delta\tau_y$ (Table 1). The value in each model grid cell in Figs. 4(a)-(e) correspond to a vector entry of $\mathbf{q} = (q_1, \dots, q_N)$ and $\mathbf{v}^* = (v_1, \dots, v_N)$, as visualized in the insets in Figs. 4(a),(b).

The Q_{net} , EPR , and τ_x vector components are not shown, but their relative contributions to the magnitude of the vectors \mathbf{q} and \mathbf{v}^* are illustrated in Figs. 4(f)-(j). The relative contribution of each forcing field F_m to \mathbf{q} is computed via

$$\|\mathbf{q}_{F_m}\|^2 = (\sigma_{\text{HT}}^{\mathbf{B}})^{-2} \sum_{i,j} \left(\frac{\partial(\text{HT}_{\text{ISR}})}{\partial F_m(i,j)} \Delta F_m \right)^2, \quad m = 1, \dots, 4, \quad (14)$$

and similarly for \mathbf{v}^* . The dark green bars in Figs. 4(f)-(j) are the area-integrated, normalized and prior-weighted sensitivities from Figs. 4(a)-(e). Put differently, the ratios in Figs. 4(f)-(j) measure the relative importance of each forcing field, F_m , for impacting HT_{ISR} and θ^* , respectively. Wind forcing (τ_x , τ_y) is more influential than buoyancy forcing (Q_{net} , EPR) for HT_{ISR} and for all observations θ^* , except for the surface observation θ^C , which is highly sensitive to local air-sea heat fluxes (Fig. 4(i), and Fig. 4(c) in Loose et al., 2020).

The yellow labels in Fig. 4 present the DPP (eq. (11)) of θ^* for HT_{ISR} , given by

$$\text{DPP}^* = \text{DPP}(\theta^*; \text{HT}_{\text{ISR}}) = (\mathbf{q} \bullet \mathbf{v}^*)^2. \quad (15)$$

The vectors \mathbf{q} and \mathbf{v}^* are composed of the sensitivities of the QoI and observations, respectively, to all control variables, capturing all possible dynamical mechanisms via which perturbations in the control variables can change the QoI and observations. Therefore, DPP^* (eq. (15)) evaluates the degree to which HT_{ISR} and θ^* share their adjustment physics. Since wind adjustments dominate HT_{ISR} and (most) observations θ^* (Figs. 4(f)-(j)), a comparison of the τ_y sensitivity map in Fig. 4(a) with the τ_y sensitivity maps in Figs. 4(b)-(e) elucidates the dynamical mechanisms and teleconnections that generate proxy potential. These mechanisms are discussed in detail in Loose et al. (2020) and only briefly reviewed in the next paragraph.

In Figs. 4(a)-(e), sensitivities emerge in two main regions: (I) in the coastal wave guide along the eastern boundary of the subtropical North Atlantic; and (II) in topographic wave guides in the northeast Atlantic and the Nordic Seas (see Figs. 5(a)-(c) for labeled regions). Wind forcing in region (I) drives a pressure adjustment mechanism (Jones et al., 2018; Loose et al., 2020) which alters both the ISR geostrophic transport and the Irminger Current (Fig. 3), causing anomalies in HT_{ISR} and in all temperature observations (Figs. 4(a)-(e)). Wind forcing in region (II) spins up an anomalous barotropic circulation around Iceland (Loose et al., 2020), which simultaneously alters the Norwegian Atlantic and East Greenland Currents (Fig. 3), causing anomalies in HT_{ISR} and in the temperature observations θ^A , θ^C , θ^D (Figs. 4(a),(b),(d),(e)). HT_{ISR} and subsurface temperature in Denmark Strait, θ^D , are the quantities that are most sensitive to the latter mechanism, driven by wind forcing in region (II). Consequently, θ^D has the largest DPP for HT_{ISR} among all observations considered (30%). The lower DPP of θ^A (19%) is explained by the strong sensitivity of information propagation because wind forcing in region (I) causes responses in HT_{ISR} and θ^A of equal sign, while wind forcing in region (II) causes responses in HT_{ISR} and θ^A of opposite sign (Figs. 4(a),(b)). The DPP of θ^B and θ^C is only 1%, since θ^B is not sensitive to wind forcing in region (II) (Fig. 5(c)), and θ^C , as a surface observation, is mostly sensitivity to local forcing (Fig. 5(d)). These characteristics lead to small sensitivity overlaps of HT_{ISR} with θ^B and θ^C .

3.2.2 Data Redundancy versus Complementarity

Next, we demonstrate how DPP generalizes when considering multiple but still noise-free observations. We begin by combining θ^A with either of the remaining temperature




Figure 5. (a)–(c): Replots of subpanels in Fig. 4, as indicated by the gray labels. The black boxes (I) and (II) contain the dominant sensitivity patterns that are the main origin of proxy potential for HT_{ISR} (see section 3.2.1). (d), (e): Modified sensitivity maps from (b), (c), as a result of extracting independent sensitivity information from the observation θ^A and θ^* , for (d) $\star = B$, (e) $\star = D$. The modified sensitivity maps are the τ_y component of $(\mathbf{v}^*)^\perp$, computed as linear combinations of \mathbf{v}^A and \mathbf{v}^* , as shown in (g), (h). (f): Three-dimensional subspace of the control space that is informed by the triple $\{\theta^A, \theta^B, \theta^D\}$. (g), (h): Planes embedded in (f), showing the two-dimensional subspaces informed by the pairs (g) $\{\theta^A, \theta^B\}$ and (h) $\{\theta^A, \theta^D\}$. The planes are spanned by (g) \mathbf{v}^A and \mathbf{v}^B , with enclosed angle $\beta = 74^\circ$, and (h) \mathbf{v}^A and \mathbf{v}^D , with enclosed angle $\delta = 30^\circ$. Orthonormalizing the pair $\{\mathbf{v}^A, \mathbf{v}^*\}$ results in (g) $(\mathbf{v}^B)^\perp$ and (h) $(\mathbf{v}^D)^\perp$.

observations θ^* , $\star = B, C, D$. In eq. (1), the resulting observing system is represented by $\mathbf{Obs} = (\theta^A, \theta^*)^T$. To tease out independent sensitivity patterns between θ^A and θ^* , we orthonormalize the vector pair $\{\mathbf{v}^A, \mathbf{v}^*\}$ (step (M.1) in Appendix A). This results in the orthonormal vector pair $\{\mathbf{v}^A, (\mathbf{v}^*)^\perp\}$, where

$$(\mathbf{v}^*)^\perp = \frac{\mathbf{v}^* - (\mathbf{v}^* \bullet \mathbf{v}^A) \mathbf{v}^A}{\|\mathbf{v}^* - (\mathbf{v}^* \bullet \mathbf{v}^A) \mathbf{v}^A\|}, \quad \star = B, C, D. \quad (16)$$

The observations θ^A and θ^B contain independent information. Indeed, θ^A is sensitive to wind forcing in both regions (I) and (II) (Fig. 5(a)), whereas θ^B is sensitive to wind forcing only in region (I) (Fig. 5(b)). Viewed in the $\{\theta^A, \theta^B\}$ -informed subspace of the control space (Fig. 5(g)), independent sensitivity information corresponds to vectors \mathbf{v}^A and \mathbf{v}^B being close to orthogonal, with an enclosed angle of $\beta = 74^\circ$. Consequently, $(\mathbf{v}^B)^\perp$ is only a slight modification of \mathbf{v}^B (Figs. 5(b),(d),(g)). In contrast to θ^B , the observation θ^D shows very similar sensitivity to wind forcing as θ^A , concentrated in both regions (I) and (II) (Figs. 5(a),(c)). Similar sensitivity information is reflected by an angle of only $\delta = 30^\circ$ between \mathbf{v}^A and \mathbf{v}^D (Fig. 5(h)). Orthonormalization of the vector pair results in a removal of sensitivity to wind forcing in regions (I) and (II) (Fig. 5(e)). The independent sensitivity information extracted from the observing systems (Fig. 5) is independent of the chosen QoI.

We now return to our QoI, and use the independent sensitivity patterns obtained in Fig. 5 to quantify the DPP of the system $\{\theta^A, \theta^*\}$ for HT_{ISR} . DPP is obtained through a projection of \mathbf{q} onto the $\{\theta^A, \theta^*\}$ -informed plane (Fig. 2(c), eq. (11)). This projection requires an orthonormal basis of the $\{\theta^A, \theta^*\}$ -informed plane, for which one may choose either $\{\mathbf{v}_1, \mathbf{v}_2\}$ or $\{\mathbf{v}^A, (\mathbf{v}^*)^\perp\}$. Here, we choose the latter, i.e., we compute

$$\text{DPP}(\theta^A, \theta^*; \text{HT}_{\text{ISR}}) = (\mathbf{q} \bullet \mathbf{v}^A)^2 + (\mathbf{q} \bullet (\mathbf{v}^*)^\perp)^2 = \text{DPP}^A + (\mathbf{q} \bullet (\mathbf{v}^*)^\perp)^2. \quad (17)$$

The first term on the right hand side is equal to 19% (cf. Fig. 4(b)) and forms the baseline value in Figs. 6(a)-(f). The second term shows the gain in DPP when adding θ^* to θ^A , and is displayed by the hatched contribution in Figs. 6(b)-(d). The generation of the hatched contribution is further explained in Figs. 6(g)-(i). The left, solid bar in each panel shows the DPP of each individual observation, $\text{DPP}^* = (\mathbf{q} \bullet \mathbf{v}^*)^2$, computed through the projection of prior-weighted and normalized sensitivity vectors (cf. Figs. 4(c)-(e)). The right, hatched bar in each panel shows $(\mathbf{q} \bullet (\mathbf{v}^*)^\perp)^2$; this projection uses the modified sensitivity patterns from Fig. 5.

The consequence of orthonormalization now becomes apparent. When combining θ^A and θ^B , data complementarity results in a DPP of $(19 + 6)\% = 25\%$ (Fig. 6(b)), exceeding the sum of $\text{DPP}^A = 19\%$ and $\text{DPP}^B = 1\%$ (Fig. 6(g)). Put differently, considering θ^A and θ^B in combination helps to extract some of the observations' sensitivity information which is lost in destructive interference when treating θ^A or θ^B in isolation, as was described in the last paragraph of section 3.2.1. Data complementarity can be viewed from yet another angle, when inspecting the position of \mathbf{q} projected onto the $\{\theta^A, \theta^B\}$ -informed subspace, denoted by \mathbf{q}_{obs} (Fig. 6(j)). \mathbf{q}_{obs} is not aligned with either of the vectors \mathbf{v}^A or \mathbf{v}^B . Therefore, a true information gain in the QoI is achieved when combining the observations θ^A and θ^B : the length of \mathbf{q}_{obs} increases, when advancing from the θ^A -informed and θ^B -informed components (orange dots, Fig. 6(j)) to the $\{\theta^A, \theta^B\}$ -informed component (radius of orange circle, Fig. 6(j)).

Adding the observations θ^D or θ^C to θ^A involves a certain degree of data redundancy, which is quantified in Figs. 6(h),(i). Proxy potential of θ^A for HT_{ISR} originates in wind forcing in regions (I),(II) (Fig. 5(a)); this sensitivity information is already contained in θ^D (Fig. 5(c), Fig. 6(k)). Consequently, $\text{DPP}(\theta^A, \theta^D; \text{HT}_{\text{ISR}})$ does not exceed $\text{DPP}^D = 30\%$ (Figs. 6(c),(h)). Similarly, the relevant sensitivity information contained in the Irminger Sea surface observation θ^C is already contained in the Irminger Sea sub-




Figure 6. (a)-(f): DPP (eqs. (17),(11)) of six different observing systems for the QoI, HT_{ISR} . Each observing system is formed by the temperature observations θ^* contained in the vector of colored labels below each subpanel. $DPP^A = 19\%$ (gray bar) forms the baseline value, and the hatched contributions show the gain in DPP when adding more observations. (g)-(i): The hatched bars (right bar in each of the three tuples) coincide with the hatched contributions in (b)-(d) above the baseline value. The height of the hatched bars is equal to $(\mathbf{q} \bullet (\mathbf{v}^*)^\perp)^2$, for (g) $\star = B$, (h) $\star = D$, (i) $\star = C$. The height of the solid bars (left bar in each of the three tuples) is equal to $DPP^* = (\mathbf{q} \bullet \mathbf{v}^*)^2$. In each of the three panels (g)-(i), an increase (decrease) in bar height from left to right indicates data complementarity (redundancy) between θ^A and θ^* . (j),(k): Replots of Figs. 5(g),(h). \mathbf{q}_{obs} (orange vector, cf. Fig. 2(c)) is the $\{\theta^A, \theta^*\}$ -informed component of \mathbf{q} , for (j) $\star = B$, (k) $\star = D$. The orange dashed circle shows the length of \mathbf{q}_{obs} , and the orange dots the projections of \mathbf{q}_{obs} onto \mathbf{v}^A and \mathbf{v}^* . (l) Same as (j),(k), but for $\star = C$.

Table 2. Observational noise ε_* , prior uncertainty $\sigma_*^{\mathbf{B}}$ (eq. (6)), sensitivity-to-noise ratio λ^* (eq. (18)), and effectiveness $\eta^* = \lambda^*/(\lambda^* + 1)$, for each observation θ^* , $\star = A, B, C, D$.

Obs	Location	ε_*	$\sigma_*^{\mathbf{B}}$	λ^*	$\lambda^*/(\lambda^* + 1)$
θ^A	Irminger Sea (subsurface)	0.1 °C	0.048 °C	0.23	19%
θ^B	Portuguese Coast (subsurface)	0.1 °C	0.059 °C	0.35	26%
θ^C	Irminger Sea (surface)	0.1 °C	0.230 °C	5.29	84%
		0.2 °C	0.230 °C	1.32	57%
		0.3 °C	0.230 °C	0.59	37%
θ^D	Denmark Strait (subsurface)	0.1 °C	0.071 °C	0.50	33%

surface observation θ^A (Fig. 6(l)). Thus, θ^C does not lead to a gain in DPP when added to θ^A (Fig. 6(d)).

Finally, we are interested in the maximum achievable DPP for HT_{ISR} , obtained by combining all four observations in our case study. Viewed within the three-dimensional subspace that is informed by the observing system $\{\theta^A, \theta^B, \theta^D\}$ (Fig. 5(f)), the $\{\theta^A, \theta^D\}$ -informed yellow plane is almost orthogonal to the $\{\theta^A, \theta^B\}$ -informed green plane (where the black plane would be exactly orthogonal to the green plane). Hence, when adding θ^D to the observing system $\{\theta^A, \theta^B\}$, the gain in DPP (green-yellow hatched, Fig. 6(e)) almost coincides with $(\mathbf{q} \bullet (\mathbf{v}^D)^\perp)^2$ (yellow hatched, Fig. 6(c)), leading to a total DPP of 35% (Fig. 6(e)). Completing the observing system by θ^C does not increase the DPP any further (Fig. 6(f)).

3.3 Noisy Observations

So far, our analysis has assumed noise-free observations. Next, we study the EPP of our observations θ^* ; this notion does account for observational noise, as encoded in the Gaussian noise matrix \mathbf{R} (eq. (1)). Recall that the EPP of θ^* for HT_{ISR} is equal to the relative uncertainty reduction in HT_{ISR} that is achieved when adding θ^* to the underlying state estimation framework (eq. (12)). Following the common assumption of uncorrelated observation errors (e.g., Forget et al., 2015), we only need to specify the diagonal entries of \mathbf{R} , i.e., the error variance ε_*^2 of each observation θ^* . We assign $\varepsilon_* = 0.1$ °C for all observations (Table 2). We also consider the impact of varying ε_C , by testing for $\varepsilon_C = 0.2$ °C and $\varepsilon_C = 0.3$ °C. The rationale for this addition is that climatological surface temperature, measured by θ^C , is more variable than climatological subsurface temperature (Locarnini et al., 2013), and can therefore be expected to be more noisy.

3.3.1 Sensitivity-To-Noise Ratio

The strength of the constraint provided by each individual observation θ^* is quantified by the eigenvalue λ^* (eq. (5)) corresponding to the θ^* -informed direction \mathbf{v}^* (Fig. 2(a)). It is given by

$$\lambda^* = \frac{(\sigma_*^{\mathbf{B}})^2}{\varepsilon_*^2} = \frac{1}{\varepsilon_*^2} \sum_{m=1}^4 \sum_{i,j} \left(\frac{\partial \theta^*}{\partial F_m(i,j)} \Delta F_m \right)^2. \quad (18)$$

λ^* describes the *sensitivity-to-noise ratio* (SensNR): it is large if either θ^* has high over-all prior-weighted sensitivity, $(\sigma_*^{\mathbf{B}})^2$, or if the observational noise ε_*^2 is small. Since surface temperature is much more sensitive to atmospheric forcing than subsurface temperature ($\sigma_C^{\mathbf{B}} \gg \sigma_*^{\mathbf{B}}$, $\star = A, B, D$, Table 2), the SensNR of θ^C is higher than that of $\theta^A, \theta^B, \theta^D$ (Fig. 7). This remains true if the noise variance for θ^C (i.e., ε_C^2) is assumed four – or even nine – times as large as that of the subsurface observations (Fig. 7).




Figure 7. The black curve is the function $\lambda \mapsto \lambda/(\lambda+1)$. The colored lines map the SensNR of θ^* (λ^* , eq. (18), circles) to the effectiveness of θ^* ($\eta^* = \lambda^*/(\lambda^* + 1)$, diamonds), cf. the values in Table 2. An observation that falls into the light gray rectangle has SensNR smaller than 1.

Note that $(\sigma_\star^B)^2$ is equal to the prior uncertainty in the observed quantity θ^* (cf. eq (8)), i.e. the uncertainty given the prior knowledge in the ocean state estimate, before taking the actual measurement. Thus, observations with SensNR smaller than 1 (here: $\theta^A, \theta^B, \theta^D$, gray rectangle in Fig. 7) are characterized by a prior uncertainty, $(\sigma_\star^B)^2$, that is smaller than their assumed observational uncertainty, ε_\star^2 .

The EPP of θ^* for HT_{ISR} is given by $\text{EPP}^* = \eta^* \cdot \text{DPP}^*$ (eq. (12)), with factor

$$\eta^* = \frac{\lambda^*}{\lambda^* + 1} < 1.$$

The factor η^* indicates what fraction of DPP^* can be retrieved and will therefore be referred to as the ‘effectiveness’ of the observation θ^* . Note that, in contrast to DPP^* , the observation’s effectiveness, η^* , is independent of the QoI under consideration. Instead, it is solely determined by the observation’s SensNR λ^* . Since the function $\lambda \mapsto \lambda/(\lambda+1)$ increases monotonically with λ (Fig. 7), observations with higher SensNRs are more effective. Therefore, the effectiveness of the surface observation θ^C is higher than that of the subsurface observations $\theta^A, \theta^B, \theta^D$ (Fig. 7). In fact, the effectiveness of $\theta^A, \theta^B, \theta^D$ is less than 50%, due to their SensNR being less than 1 (gray rectangle, Fig. 7).

In this section, we studied the SensNR, λ^* , and associated effectiveness, η^* , of each *individual* observation θ^* . In the next section, we will establish a connection between the λ^* and the set of eigenvalues $\{\lambda_i\}_{i=1}^4$, where the latter set characterizes the observing system that is *jointly* formed by $\{\theta^A, \theta^B, \theta^C, \theta^D\}$.

3.3.2 Combining Noisy Observations

We now combine all four temperature observations of our case study, while taking into account their observational noise. In eq. (1), the resulting observing system is




Figure 8. Eigen-decomposition $\{\mathbf{v}_i, \lambda_i\}_{i=1}^4$ of $\mathbf{H}_{\text{misfit}}$ (eq. (3)) for the observing system in eq. (19), with $\varepsilon_\star = 0.1^\circ\text{C}$ for $\star = A, B, D$ and $\varepsilon_C = 0.2^\circ\text{C}$. (a),(b): Orientation of the eigenvectors $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4$ (purple vectors/dots) within the (a) $\{\theta^A, \theta^C\}$ -informed, (b) $\{\theta^A, \theta^B, \theta^D\}$ -informed subspace of the control space (cf. Fig. 6(l), Fig. 5(f)). The ellipses in (a) show the contour lines of $J_{\text{misfit}}(\tilde{\mathbf{u}})$. (c)-(f): τ_y component of the four eigenvectors. Each inset reports the eigenvalue λ_i , and the associated effectiveness $\eta_i = \lambda_i / (\lambda_i + 1)$.

represented by

$$\mathbf{Obs} = (\theta^A, \theta^B, \theta^C, \theta^D)^T, \quad \mathbf{R} = \text{diag}(\varepsilon_A^2, \varepsilon_B^2, \varepsilon_C^2, \varepsilon_D^2), \quad (19)$$

where the latter denotes a diagonal 4×4 matrix, with diagonal entries equal to the noise variances ε_\star^2 , chosen as in Table 2. For the sake of brevity, we focus on the case $\varepsilon_C = 0.2^\circ\text{C}$ (cases with alternative choices for ε_C are presented in the supporting information, Fig. S.1). We compute the eigenvectors and eigenvalues, $\{\mathbf{v}_i, \lambda_i\}_{i=1}^4$, of the misfit Hessian $\mathbf{H}_{\text{misfit}}$ (eq. (3)) as described in Appendix A.

By definition, the first eigenvector \mathbf{v}_1 points in the direction of maximal curvature of $J_{\text{misfit}}(\tilde{\mathbf{u}})$. This direction is almost aligned with the θ^C -informed direction, spanned by \mathbf{v}^C (Fig. 8(a)), because the surface observation θ^C has a much higher SensNR than the remaining observations (Fig. 7). The remaining eigenvectors, $\mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4$, have little contribution from θ^C (purple dots, Fig. 8(a)), and are instead a linear combination of \mathbf{v}^\star , $\star = A, B, D$ (Fig. 8(b)). The τ_y component of \mathbf{v}_2 (Fig. 8(d)) extracts the dominant sensitivity patterns along the eastern boundary of the North Atlantic (region (I)), shared by θ^A , θ^B , and θ^D , and in the northeast Atlantic and the Nordic Seas (region (II)), shared by θ^A and θ^D (Figs. 5(a)-(c)). The τ_y component of \mathbf{v}_3 (Fig. 8(e)) is governed by sensitivities in region (I), as set by θ^B . Meanwhile, the τ_y component of \mathbf{v}_4 (Fig. 8(f)) is dominated by sensitivity dipoles local to the three observing sites (yellow dots), which emerge due to the effect of Ekman pumping.

The eigenvalues λ_i (inserted in Figs. 8(c)-(f)) are closely linked to the SensNRs λ^\star (Table 2) of all observations θ^\star involved, via the following relations (Bunch et al., 1978):

$$\sum_{i=1}^4 \lambda_i = \lambda^A + \lambda^B + \lambda^C + \lambda^D \quad \text{and} \quad \lambda_1 \geq \max\{\lambda^A, \lambda^B, \lambda^C, \lambda^D\}. \quad (20)$$

Each eigenvalue λ_i determines the effectiveness, $\eta_i = \lambda_i/(\lambda_i + 1)$, of the eigenvector \mathbf{v}_i (insets in Figs. 8(c)-(f) and diamonds in Fig. 9(b)). We consider how the effectiveness of the observing system components changes with observational noise, and inflate the observational noise covariance by a factor α , to $\alpha\mathbf{R}$. As α varies from 0 (no noise) to 1 (full noise), effectiveness decays as $\lambda/(\lambda + \alpha)$, from 100% to $\lambda/(\lambda + 1)$ (eq. (B1)). Here, λ is a placeholder for either a SensNR λ^\star (Fig. 9(a)) or an eigenvalue λ_i (Fig. 9(b)). The decay in effectiveness of the surface observation θ^C (Fig. 9(a)) as well as \mathbf{v}_1 (Fig. 9(b)) is slower than that of the remaining observations and eigenvectors.

Decay in effectiveness causes decay in proxy potential. When considering individual observations θ^\star , inflating observational noise leads to a decay in proxy potential according to $\lambda^\star/(\lambda^\star + \alpha) \cdot (\mathbf{q} \bullet \mathbf{v}^\star)^2$ (Fig. 9(c)). If instead, the full observing system is considered jointly, the decay in proxy potential is given by

$$\sum_{i=1}^4 \frac{\lambda_i}{\lambda_i + \alpha} (\mathbf{q} \bullet \mathbf{v}_i)^2, \quad (21)$$

see Fig. 9(d) (and Appendix B). The expression in equation (21) involves the eigenvectors and eigenvalues of the misfit Hessian. For noise-free observations, proxy potential is equal to DPP ($\alpha = 0$, pentagons in Figs. 9(c),(d), eq. (11)). It decays to EPP for fully inflated noise ($\alpha = 1$, squares in Figs. 9(c),(d), eq. (12)). Even though the surface observation θ^C has highest SensNR and, thus, slowest decay in effectiveness (Fig. 9(a)), its proxy potential for HT_{ISR} is lower than that of the subsurface observations θ^A and θ^D , due to its almost negligible DPP at the very outset $\alpha = 0$ (Fig. 9(c)). Through a similar argument, \mathbf{v}_1 contributes less to proxy potential than \mathbf{v}_2 and \mathbf{v}_3 (Fig. 9(d)), despite its relatively highest effectiveness (Fig. 9(b)). The insignificance of θ^C implies that proxy potential of the observing system $\{\theta^A, \theta^B, \theta^C, \theta^D\}$ for HT_{ISR} is essentially insensitive to the choice of the observation error ε_C (Figs. S1(g)-(i)).




Figure 9. (a),(b) Decay in effectiveness of (a) each individual observation θ^* and (b) the eigenvectors \mathbf{v}_i of the combined observing system (eq. (19)), as a function of α . Increasing the parameter α inflates the observational noise ($\alpha\mathbf{R}$) from no noise ($\alpha = 0$) to full noise ($\alpha = 1$). Without noise, all observations have an equal effectiveness of 100%. The colored diamonds repeat the values for $\lambda/(\lambda + 1)$ from Fig. 7 and Figs. 8(c)-(f). (c),(d) Decay in proxy potential for the QoI, HT_{ISR}, again as a function of α . Without noise, proxy potential is equal to DPP (pentagons, cf. Figs. 4(b)-(e), Fig. 6(f), eq. (11)); but decays to EPP (squares, eq. (12)) for fully inflated noise. The black dashed curve in (c) coincides with the one in (d), and shows proxy potential for all observations combined.

4 Discussion

Hessian UQ and optimal observing system design have remained underexplored computational tools in oceanography, despite their successful application by the NWP and CSE communities. In this paper, we provided dynamical insight into Hessian UQ and how to leverage this method to design ocean climate observing systems. Our results warrant some general, conceptual remarks (section 4.1), as well as discussion of specific implications for our North Atlantic case study (section 4.2). We conclude with a discussion of dimension reductions of the Hessian (section 4.3), limitations (section 4.4), and an outlook (section 4.5).

4.1 Conceptual considerations

In the context of Hessian UQ, optimality refers to a well-defined goal of the observing system, often expressed in terms of one or several QoIs to be monitored (e.g., ocean transports) or predicted (e.g., sea-ice area). Given such a QoI, we rephrased the degree of optimality of an observing system in terms of ‘proxy potential’ (eq. (12)), defined as the reduction in QoI uncertainty (on a scale of 0% to 100%) that would be achieved if the observing system was added to the ocean state estimate. We showed that proxy potential combines three main concepts: (i) the degree of shared adjustment physics between QoI and observations, measured by the projections $\mathbf{q} \bullet \mathbf{v}_i$ (Fig. 4); (ii) data redundancy versus complementarity of the distinct members of the observing system, through orthogonality of $\{\mathbf{v}_i\}$ (Figs. 5, 6); and (iii) the sensitivity-to-noise ratios (SensNR) of the observational assets, which determine the effectiveness of the observing system, through multiplication by the factors $\lambda_i/(\lambda_i + 1)$ (Fig. 9).

Concept (i) can be interpreted as the dynamical analogue of statistical correlation between QoI and observations, where proxy origins (in the space of uncertain control variables) are unambiguously identified by the adjoint model (Loose et al., 2020).

Concept (ii) is associated with the eigenvectors \mathbf{v}_i of the misfit Hessian. Given their orthogonal structure, the eigenvectors may be compared to statistical empirical orthogonal functions (EOFs), but with the following important distinction. Whereas EOF patterns are based purely on statistical evidence, the Hessian eigenvector patterns are computed through the model’s linearized dynamics. The eigenvector patterns arise as a linear combination from sensitivities of the observed quantities that are part of the observing system under investigation. They elucidate the dynamical connections between changes registered by the observing system and (local and remote) oceanic perturbations back in time that cause these observed changes (Fig. 8).

Concept (iii) is linked to the eigenvalues λ_i of the misfit Hessian. The eigenvalues quantify the strength of the constraints imposed by the observing system, and are determined by the SensNR of all observations that are part of the observing system. The SensNR of an observation can be regarded as the dynamical analogue of the statistical signal-to-noise ratio (SNR). The SensNR accounts for all possible forcing scenarios (as captured by changes in the control variables), while the SNR is based on the statistics of observed or simulated data, which samples only certain forcing occurrences. The leading eigenvectors inherit the sensitivity patterns from observations with highest SensNR (Figs. 8, S.1). The higher the observations’ SensNRs (or eigenvalues), the more efficiently the observations reduce uncertainty in the state estimate and QoIs (Fig. 9).

4.2 Implications for our case study

We found that wind anomalies in the coastal and topographic wave guides, along the eastern and northern boundaries of the North Atlantic, are the largest origins of proxy potential (Fig. 4). Such wind anomalies drive pressure adjustment mechanisms with a basin-wide response in North Atlantic circulation and temperature (Loose et al., 2020):




Figure 10. (a),(b): Replot of Figs. 4(a),(b). (c): Definition of eight ‘large regions’ covering the global ocean. (d),(e): Same τ_y sensitivities as in (a),(b), but now accumulated over the eight regions defined in (c). The maps in (d),(e) are the τ_y restrictions of \mathbf{q} and \mathbf{v}^A , if the control space is reduced a-priori from $\mathcal{O}(10^6)$ dimensions (Table 1) to only 32, via the ‘large region approach’. The yellow labels show the values for dynamical proxy potential (DPP, eq. (11)), using the full control space (left, cf. Fig. 4(b)) and the a-priori reduced control space (right). On the right, DPP evaluates to 0% due to an entire cancellation in the projection $\mathbf{q} \bullet \mathbf{v}^A$. For instance, sensitivities of equal sign in the subtropical Atlantic (region 2, (d) vs. (e)) make a strongly positive contribution to $\mathbf{q} \bullet \mathbf{v}^A$, while sensitivities of opposite sign in the Norwegian Sea (region 5, (d) vs. (e)) make a strongly negative contribution.

adjustments that are registered by QoIs and observations alike, even if the observing system is remote from the QoI. This result can be rephrased in terms of UQ (Fig. 1) as follows. North Atlantic temperature observations reduce uncertainties in various atmospheric forcing variables at various locations (via path UQ1). Among these, it is primarily uncertainty reduction in momentum fluxes along the eastern and northern boundaries of the North Atlantic which then lead to uncertainty reduction in the QoI (via path UQ2).

For a given subsurface temperature observation in the Irminger Sea, our Hessian UQ analysis reveals that a subsurface temperature observation off the Portuguese coast would provide more independent information than a subsurface temperature observation in Denmark Strait (Figs. 5, 6). This result is explained by the fact that the Irminger Sea and Denmark Strait subsurface observations have very similar adjustment physics.

Finally, our case study suggests that surface temperature observations have lower proxy potential for remote QoIs than subsurface temperature observations, despite their higher SensNR (Fig. 9). This is due to strong sensitivity of surface temperature to local air-sea fluxes, which overrides their sensitivity to the large-scale, basin-wide adjustment mechanisms that are relevant for remote QoIs (Fig. 4).

4.3 Dimension reduction

One of the main computational challenges to UQ is the curse of dimensionality (Bellman, 1957). Since the uncertain parametric model inputs (or control variables) are typically adjusted on a grid point basis of the GCM, their number is large: $\mathcal{O}(10^6)$ to $\mathcal{O}(10^8)$. The calculation of the full Hessian – a matrix with $\mathcal{O}(10^{12})$ to $\mathcal{O}(10^{16})$ elements – would require years of extensive computer resources, an intractable endeavor.

The success of Hessian UQ relies on approaches that are more computationally efficient, two of which we consider: first, an a-priori reduction, and second, a data-informed reduction of the control space dimension. In this paper we have pursued the second approach, as further discussed in the next paragraph. In contrast, Kaminski et al. (2015, 2018) follow the first approach, by aggregating and adjusting their control variables uniformly over large regions (e.g., Fig. 2 in Kaminski et al., 2015), rather than on a model grid point basis. This ‘large region approach’ reduces their control space to a total of about 150 control variables, and it is then feasible to explicitly compute the full Hessian (150^2 entries). In practice, the large region approach requires to spatially accumulate sensitivities of QoIs and observed quantities over the pre-defined large regions, as exemplified in Fig. 10. The eight regions defined in Fig. 10(c) reduce the dimension of our control space from $\mathcal{O}(10^6)$ (Table 1) to $8 \cdot 4 = 32$. However, the spatial accumulation of sensitivities implies that proxy origins and adjustment mechanisms, e.g., along the basin boundaries, are no longer resolved (Figs. 10(d),(e)) and proxy potential is artificially lost (right yellow label). Note that for other QoIs, this approach could overestimate (rather than underestimate) proxy potential and uncertainty reduction.

Because of the ad-hoc nature of a-priori control space reductions, and the difficulties it incurs (Fig. 10), we advocate the approach of data-informed reductions of the control space for the following reason. Even though the Hessian in our North Atlantic case study consists of $\mathcal{O}(10^{12})$ entries (section 3.1), the misfit Hessian is only of rank 4, equal to the number of observations involved. The four Hessian misfit eigenvectors with non-zero eigenvalues capture the Hessian’s full information. They were extracted efficiently while preserving the physical mechanism that led to uncertainty reduction. The concept of data-informed control space reduction generalizes to large, complex observing systems, e.g., mixed mooring arrays and autonomous instruments, which include thousands to millions of observations in time and space. While it becomes intractable to compute all (thousands to millions of) misfit Hessian eigenvectors, randomized numerical linear algebra for low-rank approximations can be used to extract the *leading* eigenvectors with highest eigenvalues ($M' \ll M$ in eq. (3), Bui-Thanh et al., 2012; Kalmikov & Heimbach, 2014; Liberty et al., 2007).

Moore et al. (2017) used a related technique in a regional ocean setting. They derived data-informed reduced-rank approximations of the Hessian, but with reductions sought in the observation space, rather than the control space. The two approaches are equivalent (or ‘dual’ to each other), and the implementation of the underlying variational data assimilation scheme may determine which of the two approaches is more convenient to employ. We argue that an eigen-decomposition in the control space, as suggested here, has the appeal of a straightforward dynamical attribution of proxy origins.

4.4 Limitations

Some shortcomings of the method presented should be acknowledged. Hessian UQ relies on an accurate specification of the prior and noise covariance matrices, \mathbf{B} and \mathbf{R} (eq. (1)). This is emphasized, for instance, by the fact that the relative weight of surface vs. subsurface observational noise determines the observations’ relative effectiveness, and thus the patterns that dominate the leading eigenvectors of the misfit Hessian (Fig. S.1). A second shortcoming is that the results may suffer from model dependency, a problem common to all methods for model-informed observing system design. A third limitation is that Hessian UQ makes a Gaussian approximation of the posterior probability function for the uncertain control space and the estimated ocean state space. This approximation is accurate if the linearized model provides a good representation of the ocean dynamics on the time scales investigated. The results by Loose et al. (2020) indicate that on the five-year time scale considered, nonlinearity is not a major obstacle, at least not in the non-eddy resolving model under consideration. In situations where strong nonlinearities are barriers to Gaussian approximations, Hessian UQ in combination with non-

Gaussian sampling methods have shown promise (Petra et al., 2014), but are yet to be explored in ocean and climate modeling. For instance, Stochastic Newton MCMC employs a local Gaussian approximation (given by the local inverse Hessian), which is then used as a proposal distribution for the posterior probability distribution (Petra et al., 2014).

4.5 Outlook

In our case study, we made simplifying assumptions regarding the control variables and the prior error covariance matrix (Table 1) to enable a clearer understanding of the methodology. These simplifications are readily relaxed in future work. Based on the insights gained here, we aim to compute reduced-rank approximations of the Hessian for large observing systems within the ECCO framework. Our case study highlights that the stopping criterion for truncating the eigenvalue spectrum has to be chosen carefully, because the leading Hessian eigenvectors are not always the most important ones for informing a given QoI. Indeed, eigenvectors lower down in the spectrum captured important dynamical teleconnections originating from the sensitivity of subsurface (rather than surface) observations. Future work should address the interesting question whether the abundance of surface observations (available from satellite altimetry) and their mutual complementarity (due to their local sensitivity) may be able to cover for the large-scale sensitivities of subsurface observations.

The technique presented in this paper is complementary to the more widely used OSSEs. Hessian UQ elucidates dynamical teleconnections that communicate observational constraints – via ocean currents, wave dynamics, Ekman dynamics, and geostrophy – over basin-scale distances and on monthly to interannual time scales. It provides an approach for guiding the design of observing systems that (1) maximize the information about (possibly remote) QoIs that are difficult or impossible to observe directly, and (2) are complementary to the existing observational database. We hope that Hessian UQ, in combination with OSSEs and other tools, will be more widely used for tackling the grand community challenge of co-designing a cost-effective and long-term Atlantic observing system in the coming years.

Appendix A Eigen-Decomposition of the Misfit Hessian

For an observing system with M observations, the eigen-decomposition of the misfit Hessian (eq. (3)) can be computed from the prior-weighted sensitivity vectors $\mathbf{c}_i = \mathbf{B}^{T/2} \nabla_{\mathbf{u}} \text{Obs}_i$ via the following two steps: (M.1) a QR decomposition of $\mathbf{B}^{T/2} \mathbf{A}^T$ in \mathbb{R}^N and (M.2) and an eigen-decomposition in \mathbb{R}^M .

In step (M.1), the QR decomposition of $\mathbf{B}^{T/2} \mathbf{A}^T = (\mathbf{c}_1 \mid \cdots \mid \mathbf{c}_M)$ is computed via the Gram-Schmidt process:

- $\tilde{\mathbf{w}}_1 := \mathbf{c}_1, \quad \mathbf{w}_1 = \|\tilde{\mathbf{w}}_1\|^{-1} \tilde{\mathbf{w}}_1$
- For $j = 2, \dots, M$: $\tilde{\mathbf{w}}_j = \mathbf{c}_j - \sum_{i=1}^{j-1} (\mathbf{c}_j \bullet \mathbf{w}_i) \mathbf{w}_i, \quad \mathbf{w}_j = \|\tilde{\mathbf{w}}_j\|^{-1} \tilde{\mathbf{w}}_j$

Then the $N \times M$ matrix $\mathbf{Q} := (\mathbf{w}_1 \mid \dots \mid \mathbf{w}_M)$ and the $M \times M$ matrix

$$\tilde{\mathbf{R}} = \begin{pmatrix} \|\tilde{\mathbf{w}}_1\| & \mathbf{w}_1 \bullet \mathbf{c}_2 & \mathbf{w}_1 \bullet \mathbf{c}_3 & \cdots \\ 0 & \|\tilde{\mathbf{w}}_2\| & \mathbf{w}_2 \bullet \mathbf{c}_3 & \cdots \\ 0 & 0 & \|\tilde{\mathbf{w}}_3\| & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}$$

provide the desired QR decomposition, i.e., they satisfy $\mathbf{B}^{T/2} \mathbf{A}^T = \mathbf{Q} \tilde{\mathbf{R}}$. In step (M.2), one finds an orthogonal $M \times M$ matrix \mathbf{O} and $\lambda_1 \geq \dots \geq \lambda_M \geq 0$ such that

$$\tilde{\mathbf{R}} \mathbf{R}^{-1} \tilde{\mathbf{R}}^T = \mathbf{O} \text{diag}(\lambda_1, \dots, \lambda_M) \mathbf{O}^T,$$

by means of dense matrix algebra. Combining steps (M.1) and (M.2) gives

$$\mathbf{H}_{\text{misfit}} = \mathbf{B}^{T/2} \mathbf{A}^T \mathbf{R}^{-1} \mathbf{A} \mathbf{B}^{1/2} = \mathbf{Q} \tilde{\mathbf{R}} \mathbf{R}^{-1} \tilde{\mathbf{R}}^T \mathbf{Q}^T = \mathbf{Q} \mathbf{O} \text{diag}(\lambda_1, \dots, \lambda_M) \mathbf{O}^T \mathbf{Q}^T,$$

and the i^{th} column of $\mathbf{Q} \mathbf{O}$ contains the i^{th} eigenvector of $\mathbf{H}_{\text{misfit}}$, with corresponding eigenvalue $\lambda_i \geq 0$. The eigenvectors corresponding to non-zero eigenvalues are the data-informed directions $\mathbf{v}_1, \dots, \mathbf{v}_{M'}$.

Step (M.2) is feasible, as long as the number of observations, M , is small enough to allow for dense matrix algebra in \mathbb{R}^M . For large M , one has to resort to randomized numerical linear algebra for low-rank approximations of the misfit Hessian. Such randomized algorithms continue to follow the outlined steps (M.1) and (M.2), except that the decomposition in (M.1) is substituted by an approximate, low-rank QR factorization (Halko et al., 2011; Liberty et al., 2007).

Appendix B Inflating Noise and Prior Covariances

Modifying the noise covariance matrix via $\mathbf{R} \rightarrow \alpha \mathbf{R}$ reflects a uniform deflation ($0 < \alpha < 1$) or inflation ($\alpha > 1$) of observational noise. This modification results in a reciprocal scaling of the misfit Hessian, $\mathbf{H}_{\text{misfit}} \rightarrow \mathbf{H}_{\text{misfit}}/\alpha$. Here, we substituted $\alpha \mathbf{R}$ for \mathbf{R} in eq. (3), and assume the sensitivity matrix \mathbf{A} unchanged (even though its evaluation point may change). The scaled misfit Hessian, $\mathbf{H}_{\text{misfit}}/\alpha$, has unchanged eigenvectors \mathbf{v}_i , and new eigenvalues λ_i/α . Therefore, effectiveness scales as

$$\frac{\lambda}{\lambda + 1} \rightarrow \frac{(\lambda/\alpha)}{\lambda/\alpha + 1} = \frac{\lambda}{\lambda + \alpha} \quad (\text{B1})$$

and effective proxy potential (eq. (12)) as in eq. (21).

We note that the same scaling of the misfit Hessian, $\mathbf{H}_{\text{misfit}} \rightarrow \mathbf{H}_{\text{misfit}}/\alpha$, can be achieved by modifying the prior covariance matrix via $\mathbf{B} \rightarrow \mathbf{B}/\alpha$, while keeping the noise covariance matrix unchanged. The value $\alpha = 0$ in Fig. 9 corresponds therefore either to the limit of vanishing observational noise or infinite prior uncertainty. Similarly, $\alpha = 1$ represents not only the case of unchanged \mathbf{R} and \mathbf{B} , but also the case of $\gamma \mathbf{R}$ and \mathbf{B}/γ , for any $\gamma > 0$.

Acknowledgments

We gratefully acknowledge comments by two anonymous reviewers, which helped improve the quality of this paper. We thank Gaël Forget for developing the flexible adjoint modeling tools within ECCOv4, which this work substantially benefitted from. We also thank Timothy Smith, Helen R. Pillar, Daniel N. Goldberg, Carl Wunsch, and Kerim H. Nisancioglu for useful discussions and suggestions. This study is supported by the grants NSF-OCE-1924546, NSF-OPP-1751120, NSF-OPP-1603903, and by the ECCO project via a JPL/Caltech subcontract. The ECCOv4r2 model setup used in this work can be accessed from public repositories and permanent archives (Campin et al., 2019; Forget, 2018, 2016a, 2016b). Adjoint code was generated using the TAF software tool, created and maintained by FastOpt GmbH (<http://www.fastopt.com/>).

References

- Adcroft, A., Campin, J., Dutkiewicz, S., Evangelinos, C., Ferreira, D., Forget, G., ... Molod, A. (2018). Mitgcm User Manual. *Zenodo*. doi: 1721.1/117188
- Balmaseda, M., Anderson, D., & Vidard, A. (2007). Impact of Argo on analyses of the global ocean. *Geophysical Research Letters*, 34(16). doi: 10.1029/2007GL030452
- Bellman, R. (1957). *Dynamic Programming*. Princeton University Press.

- Bennett, A. F. (1985). Array design by inverse methods. *Progress in Oceanography*, 15(2), 129–156. doi: 10.1016/0079-6611(85)90033-3
- Bui-Thanh, T., Burstedde, C., Ghattas, O., Martin, J., Stadler, G., & Wilcox, L. C. (2012). Extreme-scale UQ for Bayesian Inverse Problems Governed by PDEs. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis* (pp. 1–11).
- Bunch, J. R., Nielsen, C. P., & Sorensen, D. C. (1978, March). Rank-one modification of the symmetric eigenproblem. *Numerische Mathematik*, 31(1), 31–48. doi: 10.1007/BF01396012
- Campin, J.-M., Heimbach, P., Losch, M., Forget, G., Hill, E., Adcroft, A., ... McRae, A. T. T. (2019). *MITgcm/MITgcm: checkpoint67m*. Zenodo. doi: 10.5281/zenodo.3492298
- Chen, P. (2011). Hessian Matrix vs. Gauss-Newton Hessian Matrix. *SIAM Journal on Numerical Analysis*, 49(4), 1417–1435. doi: 10.1137/100799988
- Forget, G. (2016a). *ECCO version 4 release 2 inputs: model initialization*. Harvard Dataverse. doi: 10.7910/DVN/7XYXSF
- Forget, G. (2016b). *ECCO version 4 release 2 inputs: surface forcing fields*. Harvard Dataverse. doi: 10.7910/DVN/9WYSZF
- Forget, G. (2018). *gaelforget/eccov4: Documentation updates*. Zenodo. doi: 10.5281/zenodo.1211363
- Forget, G., Campin, J.-M., Heimbach, P., Hill, C. N., Ponte, R. M., & Wunsch, C. (2015). ECCO version 4: an integrated framework for non-linear inverse modeling and global ocean state estimation. *Geosci. Model Dev.*, 8(10), 3071–3104. doi: 10.5194/gmd-8-3071-2015
- Fujii, Y., Rémy, E., Zuo, H., Oke, P., Halliwell, G., Gasparin, F., ... Usui, N. (2019). Observing System Evaluation Based on Ocean Data Assimilation and Prediction Systems: On-Going Challenges and a Future Vision for Designing and Supporting Ocean Observational Networks. *Frontiers in Marine Science*, 6. doi: 10.3389/fmars.2019.00417
- Fujii, Y., Tsujino, H., Usui, N., Nakano, H., & Kamachi, M. (2008). Application of singular vector analysis to the Kuroshio large meander. *Journal of Geophysical Research: Oceans*, 113(C7). doi: 10.1029/2007JC004476
- Gasparin, F., Guinehut, S., Mao, C., Mirouze, I., Rémy, E., King, R. R., ... Masina, S. (2019). Requirements for an Integrated in situ Atlantic Ocean Observing System From Coordinated Observing System Simulation Experiments. *Frontiers in Marine Science*, 6. doi: 10.3389/fmars.2019.00083
- Giering, R., & Kaminski, T. (2003). Applying TAF to generate efficient derivative code of Fortran 77-95 programs. *Proc. Appl. Math. Mech.*, 2(1), 54–57. doi: 10.1145/293686.293695
- Griffa, A., Molcard, A., Raicich, F., & Rupolo, V. (2006). Assessment of the impact of TS assimilation from ARGO floats in the Mediterranean Sea. *Ocean Science*, 2(2), 237–248. doi: https://doi.org/10.5194/os-2-237-2006
- Halko, N., Martinsson, P. G., & Tropp, J. A. (2011). Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions. *SIAM Review*, 53(2), 217–288. (Publisher: Society for Industrial and Applied Mathematics) doi: 10.1137/090771806
- Halliwell, G. R., Mehari, M. F., Hénaff, M. L., Kourafalou, V. H., Androulidakis, I. S., Kang, H. S., & Atlas, R. (2017). North Atlantic Ocean OSSE system: Evaluation of operational ocean observing system components and supplemental seasonal observations for potentially improving tropical cyclone prediction in coupled systems. *Journal of Operational Oceanography*, 10(2), 154–175. doi: 10.1080/1755876X.2017.1322770
- Heimbach, P., Fukumori, I., Hill, C. N., Ponte, R. M., Stammer, D., Wunsch, C., ... Zhang, H. (2019). Putting It All Together: Adding Value to the Global Ocean and Climate Observing Systems With Complete Self-Consistent

- 794 Ocean State and Parameter Estimates. *Frontiers in Marine Science*, 6. doi:
795 10.3389/fmars.2019.00055
- 796 Heimbach, P., Wunsch, C., Ponte, R. M., Forget, G., Hill, C., & Utke, J. (2011).
797 Timescales and regions of the sensitivity of Atlantic meridional volume
798 and heat transport: Toward observing system design. *Deep Sea Research*
799 *Part II: Topical Studies in Oceanography*, 58(17-18), 1858–1879. doi:
800 10.1016/j.dsr2.2010.10.065
- 801 Hirschi, J., Baehr, J., Marotzke, J., Stark, J., Cunningham, S., & Beismann, J.-O.
802 (2003). A monitoring design for the Atlantic meridional overturning circula-
803 tion. *Geophysical Research Letters*, 30(7). doi: 10.1029/2002GL016776
- 804 Isaac, T., Petra, N., Stadler, G., & Ghattas, O. (2015). Scalable and efficient
805 algorithms for the propagation of uncertainty from data through infer-
806 ence to prediction for large-scale problems, with application to flow of the
807 Antarctic ice sheet. *Journal of Computational Physics*, 296, 348–368. doi:
808 10.1016/j.jcp.2015.04.047
- 809 Jones, D., Forget, G., Sinha, B., Josey, S., Boland, E., Meijers, A., & Shuckburgh, E.
810 (2018). Local and Remote Influences on the Heat Content of the Labrador Sea:
811 An Adjoint Sensitivity Study. *Journal of Geophysical Research: Oceans*, 0(0).
812 doi: 10.1002/2018JC013774
- 813 Kalmikov, A., & Heimbach, P. (2014). A Hessian-Based Method for Uncertainty
814 Quantification in Global Ocean State Estimation. *SIAM Journal on Scientific*
815 *Computing*, 36(5), S267–S295. doi: 10.1137/130925311
- 816 Kaminski, T., Kauker, F., Eicken, H., & Karcher, M. (2015). Exploring the utility
817 of quantitative network design in evaluating Arctic sea ice thickness sampling
818 strategies. *The Cryosphere*, 9(4), 1721–1733. doi: 10.5194/tc-9-1721-2015
- 819 Kaminski, T., Kauker, F., Toudal Pedersen, L., Voßbeck, M., Haak, H., Niederdrenk,
820 L., ... Gråbak, O. (2018). Arctic Mission Benefit Analysis: impact of sea
821 ice thickness, freeboard, and snow depth products on sea ice forecast perfor-
822 mance. *The Cryosphere*, 12(8), 2569–2594. doi: https://doi.org/10.5194/
823 tc-12-2569-2018
- 824 Köhl, A. (2005). Anomalies of Meridional Overturning: Mechanisms in the North
825 Atlantic. *Journal of Physical Oceanography*, 35(8), 1455–1472. doi: 10.1175/
826 JPO2767.1
- 827 Köhl, A., & Stammer, D. (2004). Optimal Observations for Variational Data Assim-
828 ilation. *Journal of Physical Oceanography*, 34(3), 529–542. doi: 10.1175/2513
829 .1
- 830 Leutbecher, M. (2003). A Reduced Rank Estimate of Forecast Error Vari-
831 ance Changes due to Intermittent Modifications of the Observing Net-
832 work. *Journal of the Atmospheric Sciences*, 60(5), 729–742. doi: 10.1175/
833 1520-0469(2003)060<0729:ARREOF>2.0.CO;2
- 834 Li, F., Lozier, M. S., & Johns, W. E. (2017). Calculating the Meridional
835 Volume, Heat, and Freshwater Transports from an Observing System in
836 the Subpolar North Atlantic: Observing System Simulation Experiment.
837 *Journal of Atmospheric and Oceanic Technology*, 34(7), 1483–1500. doi:
838 10.1175/JTECH-D-16-0247.1
- 839 Liberty, E., Woolfe, F., Martinsson, P.-G., Rokhlin, V., & Tygert, M. (2007). Ran-
840 domized algorithms for the low-rank approximation of matrices. *Proceedings of*
841 *the National Academy of Sciences*, 104(51), 20167–20172. doi: 10.1073/pnas
842 .0709640104
- 843 Locarnini, R., Mishonov, A., Antonov, J., Boyer, T., Garcia, H., Baranova, O., ...
844 Johnson, D. (2013, January). World Ocean Atlas 2009, Volume 1: Temper-
845 ature. *S. Levitus, Ed. NOAA Atlas NESDIS 68, U.S. Government Printing*
846 *Office, Washington, D.C., 184 pp.*
- 847 Loose, N., Heimbach, P., Pillar, H. R., & Nisancioglu, K. H. (2020). Quantifying
848 Dynamical Proxy Potential Through Shared Adjustment Physics in the North

- Atlantic. *Journal of Geophysical Research: Oceans*, 125(9), e2020JC016112. doi: 10.1029/2020JC016112
- Marshall, J., Adcroft, A., Hill, C., Perelman, L., & Heisey, C. (1997). A finite-volume, incompressible Navier Stokes model for studies of the ocean on parallel computers. *Journal of Geophysical Research: Oceans*, 102(C3), 5753–5766. doi: 10.1029/96JC02775
- Masuda, S., Awaji, T., Sugiura, N., Matthews, J. P., Toyoda, T., Kawai, Y., ... Fukasawa, M. (2010). Simulated Rapid Warming of Abyssal North Pacific Waters. *Science*, 329(5989), 319–322. doi: 10.1126/science.1188703
- Moore, A. M., Arango, H. G., Broquet, G., Edwards, C., Veneziani, M., Powell, B., ... Robinson, P. (2011). The Regional Ocean Modeling System (ROMS) 4-dimensional variational data assimilation systems: Part III - Observation impact and observation sensitivity in the California Current System. *Progress in Oceanography*, 91(1), 74–94. doi: 10.1016/j.pocean.2011.05.005
- Moore, A. M., Arango, H. G., & Edwards, C. A. (2017). Reduced-Rank Array Modes of the California Current Observing System. *Journal of Geophysical Research: Oceans*, 123(1), 452–465. doi: 10.1002/2017JC013172
- National Academies of Sciences, Engineering, and Medicine. (2017). *Sustaining Ocean Observations to Understand Future Changes in Earth’s Climate*. Washington, DC: The National Academies Press. doi: 10.17226/24919
- Perez, R. C., Garzoli, S. L., Meinen, C. S., & Matano, R. P. (2011). Geostrophic Velocity Measurement Techniques for the Meridional Overturning Circulation and Meridional Heat Transport in the South Atlantic. *Journal of Atmospheric and Oceanic Technology*, 28(11), 1504–1521. doi: 10.1175/JTECH-D-11-00058.1
- Petra, N., Martin, J., Stadler, G., & Ghattas, O. (2014). A Computational Framework for Infinite-Dimensional Bayesian Inverse Problems, Part II: Stochastic Newton MCMC with Application to Ice Sheet Flow Inverse Problems. *SIAM Journal on Scientific Computing*, 36(4), A1525–A1555. doi: 10.1137/130934805
- Pillar, H. R., Heimbach, P., Johnson, H. L., & Marshall, D. P. (2016). Dynamical Attribution of Recent Variability in Atlantic Overturning. *Journal of Climate*, 29(9), 3339–3352. doi: 10.1175/JCLI-D-15-0727.1
- Smith, T., & Heimbach, P. (2019). Atmospheric origins of variability in the South Atlantic meridional overturning circulation. *Journal of Climate*, 32(5), 1483–1500. doi: 10.1175/JCLI-D-18-0311.1
- Tarantola, A. (2005). *Inverse Problem Theory and Methods for Model Parameter Estimation*. Society for Industrial and Applied Mathematics.
- Thacker, W. C. (1989). The role of the Hessian matrix in fitting models to measurements. *Journal of Geophysical Research: Oceans*, 94(C5), 6177–6196. doi: 10.1029/JC094iC05p06177
- Zanna, L., Heimbach, P., Moore, A. M., & Tziperman, E. (2012). Upper-ocean singular vectors of the North Atlantic climate with implications for linear predictability and variability. *Quarterly Journal of the Royal Meteorological Society*, 138(663), 500–513. doi: 10.1002/qj.937
- Zhang, W. G., Wilkin, J. L., & Levin, J. C. (2010). Towards an integrated observation and modeling system in the New York Bight using variational methods. Part II: Representer-based observing strategy evaluation. *Ocean Modelling*, 35(3), 134–145. doi: 10.1016/j.ocemod.2010.06.006